

CLAIMS

1. A method of processing data in a computer system comprising at least one host and at least one content addressable storage system which stores data for the at least one host, wherein the at least one host accesses data units stored on the at least one storage system using content addresses generated based on the content of the data units, the method comprising:

- (a) maintaining at least one index that maps a content address of at least one data unit to a storage location within the at least one storage system at which the data unit is stored; and
- (b) maintaining a cache of the location index.

2. The method of claim 1, wherein the act (a) comprises an act of maintaining the location index on at least one magnetic disk, and the act (b) comprises an act of maintaining the cache in a random access memory.

3. The method of claim 1, wherein the at least one storage system includes at least one storage node having at least one storage device and at least one access node that processes access requests from the at least one host, and wherein the act (b) further comprises an act of maintaining the cache on the at least one storage node.

4. The method of claim 1, wherein the at least one storage system includes at least one storage node having at least one storage device and at least one access node that processes access requests from the at least one host, and wherein the act (b) further comprises an act of maintaining the cache on the at least one access node.

5. The method of claim 1, wherein the at least one storage system comprises a plurality of storage nodes for storing data received from the at least one host, and wherein the at least one index is stored, at least in part, on at least two of the plurality of storage nodes.

6. The method of claim 5, wherein a number of the plurality of storage nodes on which the cache of the location index is stored is less than a number of the plurality of storage nodes on which the at least one location index is stored.
7. The method of claim 6, wherein the cache of the location index is stored on a single one of the plurality of storage nodes
8. The method of claim 5, wherein the storage system comprises a plurality of access nodes and the cache of the location index is stored on at least one of the plurality of access nodes.
9. At least one computer readable medium encoded with instructions that, when executed on a computer system, perform a method of processing data, wherein the computer system comprises at least one host and at least one content addressable storage system which stores data for the at least one host, and wherein the at least one host accesses data units stored on the at least one storage system using content addresses generated based on the content of the data units, the method comprising:
 - (a) maintaining at least one index that maps a content address of at least one data unit to a storage location within the at least one storage system at which the data unit is stored; and
 - (b) maintaining a cache of the location index.
10. The at least one computer readable medium of claim 9, wherein the act (a) comprises an act of maintaining the location index on at least one magnetic disk, and the act (b) comprises an act of maintaining the cache in a random access memory.
11. The at least one computer readable medium of claim 9, wherein the at least one storage system includes at least one storage node having at least one storage device and at least one access node that processes access requests from the at least one host, and wherein the act (b) further comprises an act of maintaining the cache on the at least one storage node.

12. The at least one computer readable medium of claim 9, wherein the at least one storage system includes at least one storage node having at least one storage device and at least one access node that processes access requests from the at least one host, and wherein the act (b) further comprises an act of maintaining the cache on the at least one access node.

13. The at least one computer readable medium of claim 9, wherein the at least one storage system comprises a plurality of storage nodes for storing data received from the at least one host, and wherein the at least one index is stored, at least in part, on at least two of the plurality of storage nodes.

14. The at least one computer readable medium of claim 13, wherein a number of the plurality of storage nodes on which the cache of the location index is stored is less than a number of the plurality of storage nodes on which the at least one location index is stored.

15. The at least one computer readable medium of claim 14, wherein the cache of the location index is stored on a single one of the plurality of storage nodes.

16. The at least one computer readable medium of claim 13, wherein, the storage system comprises a plurality of access nodes and the cache of the location index is stored on at least one of the plurality of access nodes.

17. A content addressable storage system for use in a computer system, that includes the content addressable storage system and at least one host, wherein the at least one host accesses data units stored on the content addressable storage system using content addresses generated based on the content of the data units, the content addressable storage system comprising:

- at least one storage device to store data received from the at least one host; and
- at least one controller that:

maintains at least one index that maps a content address of at least one data unit to a storage location within the content addressable storage system at which the data unit is stored; and

maintains a cache of the location index.

18. The content addressable storage system of claim 17, wherein the at least one controller maintains the location index on at least one magnetic disk and maintains the cache in at least one random access memory.

19. The content addressable storage system of claim 17, further comprising at least one storage node having the at least one storage device and at least one access node that processes access requests from the at least one host, wherein the at least one controller maintains the cache on the at least one storage node.

20. The content addressable storage system of claim 17, further comprising at least one storage node having the at least one storage device and at least one access node that processes access requests from the at least one host, wherein the at least one controller maintains the cache on the at least one access node.

21. The content addressable storage system of claim 17, further comprising a plurality of storage nodes for storing data received from the at least one host, and wherein the at least one index is stored, at least in part, on at least two of the plurality of storage nodes.

22. The content addressable storage system of claim 21, wherein a number of the plurality of storage nodes on which the cache of the location index is stored is less than a number of the plurality of storage nodes on which the at least one location index is stored.

23. The content addressable storage system of claim 22, wherein the cache of the location index is stored on a single one of the plurality of storage nodes.

24. The content addressable storage system of claim 21, wherein the storage system comprises a plurality of access nodes and the cache of the location index is stored on at least one of the plurality of access nodes.